

Amendments to the Claims

Please amend Claims 1, 11 and 21. The Claim listing below will replace all prior versions of the claims in the application.

Claim Listing

1. (Currently Amended) A method ~~for configuring a distributed storage system for a failover condition, said method comprising:~~

providing a plurality of distributed object storage managers “DOSMs” for receiving requests for files;

providing at least three intelligent storage nodes directly accessible to said DOSMs over a wide area, public access network coupling the DOSMs to the intelligent storage nodes, said intelligent storage nodes accessible to said DOSMs via public access network addresses associated with the intelligent storage nodes, each intelligent storage node including a processor core and a plurality of storage devices;

storing at least one file in a first intelligent storage node accessed via a DOSM over said network;

storing a duplicate of said file in a second intelligent storage node accessed via said network;

in the event of a failure of said first intelligent storage node resulting in a failover condition rendering said first intelligent storage node unavailable, upon receiving a request for said file by a DOSM, identifying by said DOSM that said second intelligent storage node stores said duplicate of said file, redirecting said file request, via said network, to said second intelligent storage node and indicating a location determined at said DOSM for said file in said second intelligent storage node; and

accessing, via said network, said file stored in said second intelligent storage node in response to said file request or a subsequent file request.

2. (Previously Presented) The method as set forth in claim 1, wherein:

storing at least one file in a first intelligent storage node accessed via a network comprises accessing said first intelligent storage node via a first network address;

storing a duplicate of said file in a second intelligent storage node accessed via said network comprises accessing said second intelligent storage node via a second network address; and

determining a location for said file in said second intelligent storage node comprises generating a mapping between said first network address and said second network address.

3. (Original) The method as set forth in claim 2, wherein said first network address and said second network address comprise internet protocol (“IP”) network addresses and differ only in a subnet portion of said IP network addresses.

4. (Previously Presented) The method as set forth in claim 1, wherein:

storing at least one file in a first intelligent storage node comprises storing said file in said first intelligent storage node located in a first storage center; and

storing a duplicate of said file in a second intelligent storage node comprises storing said file in said second intelligent storage node located in a second storage center, geographically distant from said first storage center.

5. (Previously Presented) The method as set forth in claim 4, further comprising:

storing a plurality of files in a plurality of intelligent storage nodes in said first storage center;

storing duplicates of said files in a plurality of intelligent storage nodes in said second storage center, so as to provide a one to one mapping between said intelligent storage nodes in said first storage center and said intelligent storage nodes in said second storage center.

6. (Previously Presented) The method as set forth in claim 1, wherein:

storing at least one file in a first intelligent storage node comprises storing said file in said first intelligent storage node located in a first storage center; and
storing a duplicate of said file in a second intelligent storage node
comprises storing said file in said second intelligent storage node located in said first storage center.

7. (Previously Presented) The method as set forth in claim 1, wherein:
storing at least one file in a first intelligent storage node comprises:
storing said file in a first storage center comprising a plurality of intelligent storage nodes; and
storing said duplicate of said file in a second storage center, geographically distant from said first storage center; and
determining a location for said file in said second intelligent storage node comprises searching for said file in said second storage after entering said failover condition.
8. (Previously Presented) The method as set forth in claim 7, wherein identifying further comprises searching for said file using a point-to-point protocol between said distributed object storage manager (DOSM) and said second intelligent storage node.
9. (Previously Presented) The method as set forth in claim 7, wherein identifying further comprises searching for said file using a multi-cast protocol.
10. (Previously Presented) The method as set forth in claim 7, identifying further comprises searching for said file using a point-to-point protocol between a said distributed object storage manager (DOSM) and an intelligent storage node.
11. (Currently Amended) A distributed storage system comprising:
a wide area, public access network;

at least three intelligent storage nodes, each intelligent storage node including a processor core and a plurality of storage devices;

a first intelligent storage node having multiple storage devices, accessed via said network for storing at least one file;

a second intelligent storage node having multiple storage devices, accessed via said network with a public access network address, for storing a duplicate of said file; and

a plurality of distributed object storage managers (DOSMs) remote from said storage nodes, one of said DOSMs for directly accessing, via an interconnection over said network, said file stored in said first intelligent storage node in response to a file request, in the event of a failure of said first intelligent storage resulting in a failover condition rendering said first intelligent storage device unavailable, upon receiving a request for said file by a DOSM, for identifying by said DOSM that said second intelligent storage node stores said duplicate of said file, for redirecting said file request to said second intelligent storage node, via said network, indicating a location for said file in one of said intelligent storage nodes and for directly accessing, via said interconnection over said network, said file stored in said second intelligent storage node in response to said file request or a subsequent file request.

12. (Original) The distributed storage system as set forth in claim 11, wherein said DOSM further for accessing said first intelligent storage node via a first network address and accessing said second intelligent storage node via a second network address, said DOSM further comprising a mapping between said first network address and said network address.
13. (Original) The distributed storage system as set forth in claim 12, wherein said first network address and said second network address comprise internet protocol (“IP”) network addresses and differ only in a subnet portion of said IP network address.

14. (Original) The distributed storage system as set forth in claim 11, wherein:
 - a first storage center comprising said first intelligent storage node; and
 - a second storage center, geographically distant from said first storage center comprising said second intelligent storage node.
15. (Original) The distributed storage system as set forth in claim 14, wherein:
 - said first storage center comprises a plurality of files stored in a plurality of intelligent storage nodes; and
 - said second storage center comprises duplicates of said files stored in a plurality of intelligent storage nodes, so as to provide a one to one mapping between said intelligent storage nodes in said first storage center and said intelligent storage nodes in said second storage center.
16. (Original) The distributed storage system as set forth in claim 11, wherein said first and second intelligent storage nodes reside in a single storage center.
17. (Original) The distributed storage system as set forth in claim 11, further comprising:
 - a first storage center comprising a plurality of intelligent storage nodes for storing said file;
 - a second storage center, geographically distant from said first storage center, for storing said duplicate of said file; and
 - said DOSM for searching for said file in said second storage center after entering said failover condition.
18. (Previously Presented) The distributed storage system as set forth in claim 17, wherein said DOSM further comprises processes for searching for said file in said first storage center if said file is not located in said second storage center.

19. (Previously Presented) The distributed storage system as set forth in claim 17, wherein identifying uses a multi-cast protocol for communicating among said DOSM and said intelligent storage nodes.
20. (Previously Presented) The distributed storage system as set forth in claim 17, wherein identifying uses a point-to-point protocol for communicating between said DOSM and one of said intelligent storage nodes.
21. (Currently Amended) A distributed virtual file system comprising:
 - a wide area, public access network;
 - a first directory, remote from a requesting client and from an associated intelligent storage node, accessed via said network, for storing file system information associated with said intelligent storage node having multiple storage devices and a processor core;
 - a second directory, accessed via said network, for storing a duplicate of said file system information; and
 - at least one distributed directory manager (DDM) for directly accessing, via said network, said file system information stored in said first directory in response to a file system request, and for redirecting, in the event of a failure of said first directory resulting in a failover condition rendering said first directory unavailable, a file request from said first directory to said second directory, via said network, indicating a location for said file system information in one of said second directory, and for directly accessing, via said network, said file system information stored in said second directory in response to said redirected file request.
22. (Original) The distributed virtual file system as set forth in claim 21, wherein said DDM further for accessing said first directory via a first network address and accessing said second directory via a second network address, said DDM further comprising a mapping between said first network address and said second network address.

23. (Original) The distributed virtual file system as set forth in claim 22, wherein said first network address and said second network address comprise internet protocol (“IP”) network addresses and differ only in a subnet portion of said IP network address.
24. (Original) The distributed virtual file system as set forth in claim 21, wherein:
 - a first storage center comprising said first directory; and
 - a second storage center, geographically distant from said first storage center comprising said second directory.
25. (Original) The distributed virtual file system as set forth in claim 24, wherein:
 - said first storage center comprises file system information stored in a plurality of directories; and
 - said second storage center comprises a duplicate of said file system information stored in a plurality of directories, so as to provide a one to one mapping between said directories in said first storage center and said directories in said second storage center.
26. (Original) The distributed virtual file system as set forth in claim 21, wherein said first and second directories reside in a single storage center.
27. (Original) The distributed virtual file system as set forth in claim 21, further comprising:
 - a first storage center comprising a plurality of directories for storing said file system information;
 - a second storage center, geographically distant from said first storage center, for storing a duplicate of said file system information; and
 - said DDM for searching for said file system information in said second storage center after entering said failover condition.

28. (Previously Presented) The distributed virtual file system as set forth in claim 27, wherein said DDM further comprises processes for searching for said file system information in said first storage center if said file system information is not located in said second storage center.
29. (Previously Presented) The distributed virtual file system as set forth in claim 27, wherein a multi-cast protocol is used for communicating among said DDM and said directories.
30. (Previously Presented) The distributed virtual file system as set forth in claim 27, wherein a point-to-point protocol is used for communicating between said DDM and one of said directories.
31. (Previously Presented) The distributed storage system as set forth in claim 11, wherein the processing core communicates with the storage devices using the ISA protocol.
32. (Previously Presented) The distributed storage system as set forth in claim 11, wherein the processing core communicates with the storage devices using the SCSI protocol.